

(12) UK Patent Application (19) GB (11) 2 230 370 A (13)

(43) Date of A publication 17.10.1990

(21) Application No 9007067.3

(22) Date of filing 29.03.1990

(30) Priority data

(31) 8908205

(32) 12.04.1989

(33) GB

(71) Applicant

Smiths Industries Public Limited Company

(Incorporated in the United Kingdom)

765 Finchley Road, London, NW11 8DS,
United Kingdom

(72) Inventor

Ian Bickerton

(74) Agent and/or Address for Service

J M Flint

765 Finchley Road, London, NW11 8DS,
United Kingdom

(51) INT CL^a

G10L 5/06

(52) UK CL (Edition K)

G4R RHB R1F R10A R10E R10X R11D R11E R11F

R12F R3A R3C R3E R3G R3X R8F R9B R9C

U1S S1288 S1839 S1840 S2409

(56) Documents cited

EP 0242743 A1

(58) Field of search


UK CL (Edition K) G4R RHB RRL RRP

INT CL^a G10L


Online databases: WPI, INSPEC

(54) Speech recognition

(57) Speech recognition is carried out by performing a first analysis of a speech signal using a Hidden Semi Markov Model and an asymmetric time warping algorithm 16, 17, 18. A second analysis is also performed using Multi-Layer Perceptron techniques in conjunction with a neural net 20. The first analysis is used by the second to identify word boundaries. Where the first analysis provides an indication of the word spoken above a certain level of confidence, an output representative of the word spoken may be generated solely in response to the first analysis, the second analysis being utilised when the level of confidence falls. The output controls 4, a function of an aircraft and provides feedback, 3, to the speaker of the words spoken.



GB 2 230 370 A



SPEECH RECOGNITION APPARATUS AND METHODS

This invention relates to speech recognition apparatus and methods.

In complex equipment having multiple functions it can be useful to be able to control the equipment by spoken commands. This is also useful where the user's hands are occupied with other tasks or where the user is disabled and is unable to use his hands to operate conventional mechanical switches and controls.

The problem with equipment controlled by speech is that speech recognition can be unreliable, especially where the voice of the speaker is altered by environmental factors, such as vibration. This can lead to failure to operate or, worse still, to incorrect operation.

Various techniques are used for speech recognition. One technique involves the use of Markov models which are useful because they readily enable the boundaries between words in continuous speech to be identified. In noisy environments or where speech is degraded by stress on the speaker, Markov model techniques may not provide sufficiently reliable identification of the words spoken. Considerable effort has been made recently to improve the performance of such techniques by

noise compensation, compensation, syntax selection and other methods.

An alternative technique which has been proposed for speech recognition employs neural nets. These neural net techniques are capable of identifying individual words to high accuracy even when speech is badly degraded. They are, however, not suited to the recognition of continuous speech because they are not capable of accurately identifying word boundaries.

It is an object of the present invention to provide improved speech recognition apparatus and methods.

According to one aspect of the present invention there is provided a method of speech recognition comprising the steps of performing a first analysis of a speech signal to identify boundaries between different words and to provide a first indication of the words spoken by comparison with a stored vocabulary, performing a second analysis of the speech signal utilising neural net techniques and word boundary identification from the first analysis to provide a second indication of the words spoken, and providing an output signal representative of the words spoken from at least said second indication.

The first analysis may be performed using a Markov model which may be a Hidden Semi Markov model. The vocabulary may contain dynamic time warping templates and the first analysis may be performed using an asymmetric dynamic time warping algorithm.

The first analysis is preferably performed utilising a plurality of different algorithms, each algorithm providing a signal indicative of the word in the vocabulary store closest to the speech signal together with an indication of the confidence that the indicated word is the word spoken, a comparison being made between the signals provided by the different algorithms. Where the first indication of the words spoken is provided with a measure of confidence, the output signal may be provided solely in response to the first indication when the measure of confidence is greater than a predetermined value.

The second analysis may be performed using a multi-layer perceptron technique in conjunction with a neural net.

The output signal may be utilised to provide feedback to the speaker of the words spoken and may be utilised to control a function of an aircraft.

According to another aspect of the present invention there is provided apparatus for carrying out a method according to the above one aspect of the present invention.

According to a further aspect of the present invention there is provided speech recognition apparatus including store means containing speech information about a vocabulary of words that can be recognised, means for performing a first analysis of a speech signal to identify boundaries between different words and to compare the speech signal with the stored vocabulary to provide a first indication of the words spoken, means for performing a second analysis of the speech signal utilising neural net techniques and word boundary identification from said first analysis to provide a second indication of the words spoken, and means for providing an output signal representative of the words spoken from at least the second indication.

The speech signal may be derived from a microphone. The apparatus may include a noise marking unit which performs a noise marking algorithm on the speech signals. The apparatus may include a syntax unit which performs syntax restriction on the stored vocabulary in accordance with the syntax of previously identified words.

Speech recognition apparatus and its method of operation in accordance with the present invention will now be described, by way of example, with reference to the accompanying drawing which shows the apparatus schematically.

The speech recognition apparatus is indicated generally by the numeral 1 and receives speech input signals from a microphone 2 which may for example be mounted in the oxygen mask of an aircraft pilot. Output signals representative of identified words are supplied by the apparatus 1 to a feedback device 3 and to a utilisation device 4. The feedback device 3 may be a visual display or an audible device arranged to inform the speaker of the words as identified by the apparatus 1. The utilisation device 4 may be arranged to control a function of the aircraft equipment in response to a spoken command recognised by the utilisation device from the output signals of the apparatus.

Signals from the microphone 2 are supplied to a pre-amplifier 10 which includes a pre-emphasis stage 11 that produces a flat long-term average speech spectrum to ensure that all the frequency channel outputs occupy a similar dynamic range, the characteristic being nominally flat up to 1 kHz. A switch 12 can be set to give either a 3 or 6 dB/octave lift at higher frequencies. The pre-amplifier 10 also includes an anti-aliasing filter 21 in the form of an 8th order Butterworth low-pass filter with a -3dB cut-off frequency set at 4 kHz.

The output from the pre-amplifier 10 is fed via an analogue-to-digital converter 13 to a digital filterbank 14. The filterbank 14 has nineteen channels implemented as assembly software in a TMS32010 microprocessor and is based on the JSRU Channel Vocoder described by Holmes, J.N in IEE Proc., Vol 127, Pt.F, No.1, Feb 1980. The filterbank 14 has uneven channel spacing corresponding approximately with the critical bands of auditory perception in the range 250-4000Hz. The responses of adjacent channels cross at approximately 3dB below their peak. At the centre of a channel the attenuation of a neighbouring channel is approximately 11dB.

Signals from the filterbank 14 are supplied to an integration and noise marking unit 15 which incorporates a noise marking algorithm of the kind described by J.S. Bridle et al. 'A noise compensating spectrum distance measure applied to automatic speech recognition.' Proc. Inst. Acoust., Windemere, Nov. 1984'. Adaptive noise cancellation techniques to reduce periodic noise may be implemented by the unit 15 which can be useful in reducing, for example, periodic helicopter noise.

The output of the noise marking unit 15 is supplied to a pattern matching unit 16 which performs the various pattern matching algorithms. The pattern matching unit 16 is connected with a vocabulary store

17 which contains Dynamic Time Warping (DTW) templates and Markov models of each word in the vocabulary.

The DTW templates can be created using either single pass, time-aligned averaging or embedded training techniques. The template represents frequency against time and spectral energy.

The Markov models are derived during training of the apparatus from many utterances of the same word, spectral and temporal variation being captured with a stochastic model. The Markov model is made up of a number of discrete states, each state comprising a pair of spectral and variance frames. The spectral frame contains nineteen values covering the frequency range from 120 Hz to 4 kHz; the variance frame contains the variance information associated with each spectral vector/feature in the form of state mean duration and standard deviation information.

The individual utterances during training are analysed to classify stationary phonetic states and their spectral transitions. The model parameters are estimated with an iterative process using the Viterbi re-estimation algorithm as described by Russell, M.J. and Moore, R.H. 'Explicit modelling of state occupancy in hidden Markov Models for automatic speech recognition',

Proc IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Tampa, 26 - 29 March 1985. The final word model contains the natural spoken word variability, both temporal and inflection.

Intermediate the store 17 and the pattern matching unit 16 is a syntax unit 18 which performs conventional syntax restriction on the stored vocabulary with which the speech signal is compared, according to the syntax of previously identified words.

The pattern matching unit 16 is also connected with Neural Net unit 20. The Neural Net unit 20 incorporates a Multi-Layer Perceptron (MLP) such as described by Peeling, S.M. and Moore, R.H. 'Experiments in isolated digit recognition using the multi-layer perceptron' RSRE Memorandum No. 4073, 1987.

The MLP has the property of being able to recognise incomplete patterns such as might occur where high background noise masks low energy fricative speech. The MLP is implemented in the manner described by Rumelhart, D.E. et al. 'Learning internal representations by error back propagation' Institute for Cognitive Science, UCSD, ICS Report 8506, September 1985.

The pattern matching unit 16 employs three different algorithms to select the best match between the spoken word and the words in the vocabulary.

One is an asymmetric DTW algorithm of the kind described by Bridle, J.S. 'Stochastic models and template matching: some important relationships between two apparently different techniques for automatic speech recognition' Proc. Inst. of Acoustics, Windemere, Nov. 1984 and by Bridle, J.S. et al 'Continuous connected word recognition using whole word templates'. The Radio and Electronic Engineer, Vol. 53, No. 4, April 1983. This is an efficient single pass process which is particularly suited for real-time speech recognition. The algorithm works effectively with noise compensation techniques implemented by the unit 15.

A second algorithm employs Hidden Semi Markov Model (HSMM) techniques in which the Markov Models contained within the vocabulary store 17 described above are compared with the spoken word signals. The additional information in the Markov Models about temporal and inflection variation in the spoken words enhances recognition performance during pattern matching. In practice, the DTW and HSMM algorithms are integrated with one another. The integrated DTW and HSMM techniques are capable of identifying boundaries between adjacent

words in continuous speech.

The third algorithm employs MLP techniques in conjunction with the Neural Net 20. The MLP is controlled by the DTW/HSMM algorithm, the MLP having a variable window of view onto a speech buffer (not shown) within the pattern matching unit 16, the size and position of this window being determined by the DTW/HSMM algorithm. In this way, the HSMM algorithm is used by the MLP to identify the word boundaries or end points and the spectral time segments or word candidates can then be processed by the MLP. Each algorithm provides a signal indicative of its explanations of the speech signal such as by indicating the word in the vocabulary store identified by the algorithm most closely with the speech, together with a confidence measure. A list of several words may be produced by each algorithm with their associated confidence measures. Higher level software within the unit 16 compares the independent results achieved by each algorithm and produces an output to the feedback device 3 and utilisation device 4 based on these results after any weighting.

In this way, the apparatus of the present invention enables Neural Net techniques to be used in the recognition of natural, continuous speech which has not previously been possible. One of the advantages of the apparatus and methods of the present invention is that it can have a short response time and provide rapid feedback to the speaker. This is particularly important in aircraft applications.

It will be appreciated that alternative algorithms may be used, it only being necessary to provide one algorithm capable of identifying word boundaries in conjunction with a second algorithm employing Neural Net techniques.

The Neural Net algorithm need not be used for every word. In some apparatus it may be arranged that the Markov algorithm alone provides the output for as long as its measure of confidence is above a certain level. When a difficult word is spoken, or spoken indistinctly or with high background noise, the measure of confidence will fall and the apparatus consults the Neural Net algorithm for an independent opinion.

It will be appreciated that the functions carried out by the units described could be carried out by programming of one or more computers and need not be performed by the discrete units referred to above.

The apparatus may be used for many applications but is especially suited for use in high noise environments, such as for control of machinery and vehicles, especially fixed-wing and rotary-wing aircraft.

CLAIMS

1. A method of speech recognition comprising the steps of performing a first analysis of a speech signal to identify boundaries between different words and to provide a first indication of the words spoken by comparison with a stored vocabulary, performing a second analysis of the speech signal utilising neural net techniques and word boundary identification from the first analysis to provide a second indication of the words spoken, and providing an output signal representative of the words spoken from at least said second indication.
2. A method according to Claim 1, wherein the first analysis is performed using a Markov model.
3. A method according to Claim 2, wherein the Markov model is a Hidden Semi Markov Model.
4. A method according to any one of the preceding claims, wherein the vocabulary contains dynamic time warping templates.

5. A method according to Claim 4, wherein the first analysis is performed using an asymmetric dynamic time warping algorithm.
6. A method according to any one of the preceding claims, wherein the first analysis is performed utilising a plurality of different algorithms, wherein each algorithm provides a signal indicative of the word in the vocabulary store closest to the speech signal together with an indication of the confidence that the indicated word is the word spoken, and wherein a comparison is made between the signals provided by the different algorithms.
7. A method according to any one of the preceding claims, wherein the said first indication of the words spoken is provided with a measure of confidence, and wherein the said output signal is provided solely in response to said first indication when the measure of confidence is greater than a predetermined value.

8. A method according to any one of the preceding claims, wherein the second analysis is performed using a multi-layer perceptron technique in conjunction with a neural net.
9. A method according to any one of the preceding claims, wherein the said output signal is utilised to provide feedback to the speaker of the words spoken.
10. A method according to any one of the preceding claims, wherein the said output signal is utilised to control a function of an aircraft.
11. A method substantially as hereinbefore described with reference to the accompanying drawing.
12. Apparatus for carrying out a method according to any one of the preceding claims.
13. Speech recognition apparatus including store means containing speech information about a vocabulary of words that can be recognised, means for performing a first analysis of a speech signal to identify boundaries between different words and to compare the speech signal with the stored vocabulary to provide a first indication of the

words spoken, means for performing a second analysis of the speech signal utilising neural net techniques and word boundary identification from said first analysis to provide a second indication of the words spoken, and means for providing an output signal representative of the words spoken from at least the second indication.

14. Apparatus according to Claim 13, wherein the means for performing the first analysis uses a Markov model.
15. Apparatus according to Claim 14, wherein the Markov model is a Hidden Semi Markov model.
16. Apparatus according to any one of Claims 13 to 15, wherein the vocabulary contains dynamic time warping templates.
17. Apparatus according to Claim 16, wherein the first analysis is performed using an asymmetric dynamic time warping algorithm.
18. Apparatus according to any one of the Claims 13 to 17, wherein the first analysis is performed utilising a plurality of different algorithms, wherein each algorithm provides a signal

indicative of the word in the vocabulary store closest to the speech signal together with an indication of the confidence that the indicated word is the word spoken, and wherein the apparatus includes means for comparing the signals provided by the different algorithms.

19. Apparatus according to any one of Claims 13 to 18, wherein the said first indication of the words spoken is provided with a measure of confidence, and wherein the said output signal is provided solely in response to said first indication when the measure of confidence is greater than a predetermined value.
20. Apparatus according to any one of Claims 13 to 19, wherein the apparatus performs the second analysis using a multi-layer perceptron technique in conjunction with a neural net.
21. Apparatus according to any one of Claims 13 to 20, including feedback means arranged to provide feedback to the speaker of the words spoken.

22. Apparatus according to any one of Claims 13 to 21 including utilisation means for controlling a function of an aircraft, and wherein the output signal is provided to the utilisation means.
23. Apparatus according to any one of Claims 13 to 22, wherein the speech signal is derived from a microphone.
24. Apparatus according to any one of Claims 13 to 23, wherein the apparatus includes a noise marking unit which performs a noise marking algorithm on the speech signals.
25. Apparatus according to any one of Claims 13 to 24, wherein the apparatus includes a syntax unit which performs syntax restriction on the stored vocabulary in accordance with the syntax of previously identified words.
26. Apparatus substantially as hereinbefore described with reference to the accompanying drawing.
27. Any novel feature or combination of features as hereinbefore described.